growth and performance. For the working trader, P&L is the bottom line. It is important to maintain clean, well-categorized, and accurate P&L records, and to have your own set of records to double-check your broker or clearing firm's accounting. Errors happen. Positions may be missing, may be sized incorrectly, or may end up in the wrong account. Corporate actions, buy-ins, dividends, rolls in futures, interest rates in forex accounts, currency conversions for assets held on foreign exchanges, and options assignments can all cause issues, so develop an obsessive habit of checking statements and records after any of these events. In many cases the customer has some liability in this type of error; treat your trading as a business and manage every aspect of it with discipline.

Though many brokers produce records that are adequate for most accounting and tax purposes, more resolution is usually needed. One of the most important tasks in record keeping is categorizing all trades and executions according to the setup or pattern that generated the trade. For instance, it is useful to be able to separate all pullback trades or all failure test trades, and then to do statistical analysis on that particular set of trades. Again, think carefully when you develop this system, because it is important to strike a balance between ease of use and sufficient detail. If you do not include enough detail, your analysis will not tell you anything useful about your performance. If you include too much detail, you may create an onerous task and eventually stop updating these records every day. In general, most traders might find it useful to have three to five main categories of trades corresponding to specific setups and patterns they trade. Each category could optionally have more information regarding the specific entry or exit technique used to allow for cross-referencing and deeper analysis. Excel is adequate for keeping these records, but many traders will find a database program better suited to the task.

It is very important that you create a process that is meaningful for your style and frequency of trading, and it is also important to get it right from the beginning. If, for instance, you change the way you categorize your trades, you may have to go back through several hundred or thousands of trades and recategorize them all. Not only is this a waste of time, but it compromises the integrity of the records; you may be inclined to push certain trades into certain categories in a subtle effort to skew the performance metrics on some classes of trades. In many cases it is possible to designate a trade several ways: Was that a pullback trade? Did you buy at support? Did you buy a failure test against support? Was it a selling climax trade on a lower time frame? What matters is what you thought it was *at the time you made the trade*, not some decisions you may make weeks or months afterward, so beware of making revisions to these designations after the fact.

STATISTICAL ANALYSIS OF TRADING RESULTS

There are so many subjective elements in discretionary trading, but your P&L is not one of them—you are either making or losing money. Correct analysis of your P&L can reveal hidden strengths and weaknesses. You may find you trade certain market environments

or asset classes better than others, or that you are trading a specific pattern well but giving back all of the P&L in other, less successful trades. In addition, tracking performance over a period of time can give early warning of problems, can highlight style drifts, and can tie performance to certain market environments. For traders at all stages of development, the main questions to ask of this analysis are:

- Are you making money? Does your P&L show that you have an edge? Are your results better than what someone could have achieved flipping a coin or throwing a dart?
- Is your performance dependent on one class of setup more than others? Could you improve performance by eliminating certain setups from your trading plan?
- Is there a market, group of markets, or asset class that you trade better than others?
- How is your performance changing compared to other time periods?

These questions can be answered through some fairly simple procedures. Many brokers will supply statements that will answer some of these questions, but others will need to be addressed through your own work and records. Once these basic questions have been answered, the analysis and review can be extended with other questions, such as:

- Are you repeating certain mistakes often enough that eliminating them could have a very positive effect on your overall P&L?
- Are you achieving superior risk-adjusted performance?
- Are there other things that you are doing that are hurting your numbers?

This next section will walk through a theoretical analysis to show the basic concepts and math used. (The set of trades is available from the author's web site at www.adamhgrimes.com/ and in Appendix C.) If this type of work is new to you, you should repeat the analysis yourself and compare your results to those in this section.

Quantifying Your Edge

Table 12.1 shows summary statistics for an active swing trader trading four separate systems over three months in a $100,000 account. Number of trades ($N=$) and cumulative net P&L (Sum) are given, as well as means for all trades and both winning and losing

TABLE 12.1 Summary Statistics for Net P&L

System	*N*=	Sum	Mean	StDev	AvgWin	AvgLoss	Win%	*p*=
A	29	(2,531)	(87)	616	456	(528)	44.8%	0.226
B	17	5,109	301	515	557	(314)	70.6%	0.014
C	16	1,109	69	653	361	(1,194)	81.3%	0.339
D	6	(529)	(88)	479	468	(366)	33.3%	0.336
All	68	3,158	46	600	456	(538)	58.8%	0.263

trades separately. The percentage of winning trades is given as Win% in the table. In this example, losing trades are 1 − Win% for all categories, but in other cases there may be breakeven trades to consider. Standard deviation is calculated as an intermediate step to deriving the p value, which is a significance test (one-tailed t-test) for the mean P&L being >0.

First, consider the bottom line of the table, which gives summary statistics for all trades combined. Over this period, this trader made $3,158 on starting capital of $100,000. The first question to ask is: Is this good performance? Even a simple question like this may not have a simple answer. On one hand, $3,158 is not a lot of money, but if this is performance over one quarter, a 3.16 percent gain is perhaps not bad. Assuming this trader can repeat this performance regularly, 3.16 percent quarterly compounds to a respectable 13.2 percent annually: $(1 + 0.0316)^4 - 1$.

However, it is not so simple: The standard deviation of the trades is a staggering $600 compared to a mean trade of $46. This is not good, and suggests that the positive performance could have been nothing more than luck. The p value of 0.263 basically says that, given this set of trades, there is a 26.3 percent chance of seeing results at least this extreme due to random chance even if the trader did not actually have an edge. Figure 12.1, a running total of the individual trades, provides another check on consistency and edge. (Note that this figure does not mirror the daily account P&L that is marked to market; this graph simply sums each trade's closed P&L. Graphs like this can be ordered either by trade entry or by exit date.) Consider this P&L line and ask yourself how hard it is to believe that it ended above zero just due to chance. In this case, there seems to be little consistency and the line wanders about as far above as below the zero

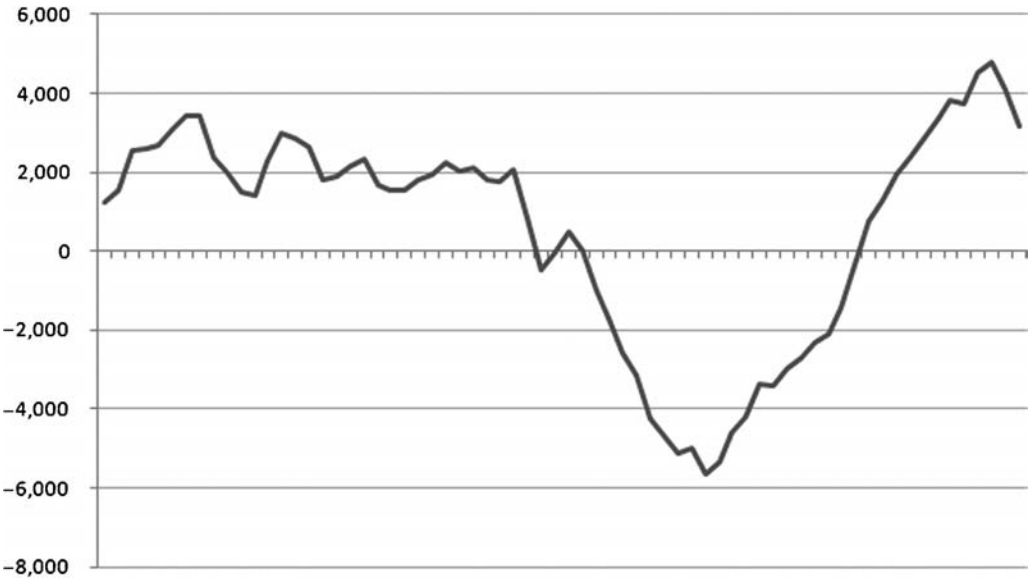

FIGURE 12.1 Cumulative P&L by Trade

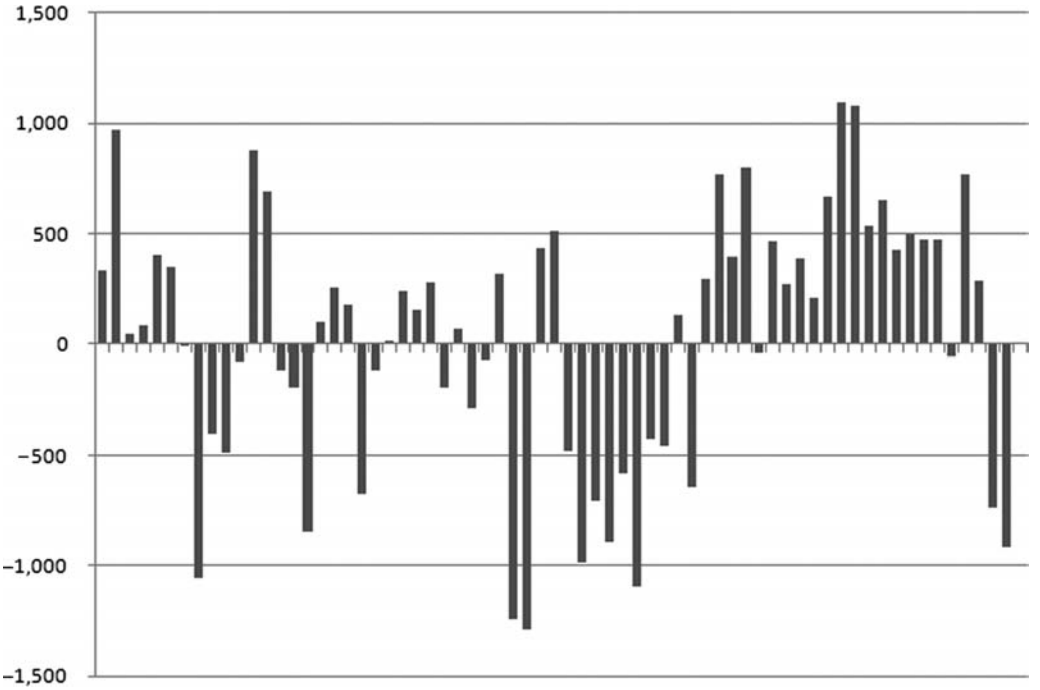

FIGURE 12.2 Trade-by-Trade Net P&L

line. This is not what we would expect to see if a trader had a strong edge. Though this is not an actual statistical test, visual inspection of a time series chart is a good companion to deeper analyses.

Yet another check comes from plotting the P&L of each individual trade (not cumulative) as in Figure 12.2. One of the most important aspects of this graph is the long runs of both winning and losing trades. There seem to be roughly three separate sections to this trader's performance: The first half of the graph shows winning and losing trades approximately equally matched, and then there is a potentially nasty series of losing trades as two large losses, followed by two small wins, lead to eight losing trades in a row. However, this period is followed by a nearly unbroken series of wins on the last third of the graph. There are questions that must be asked here that cannot be answered by this simple analysis: Why did this happen? Does the trader's style favor certain market environments over others, or is there an element of performance psychology here? Is the trader himself responsible for these streaks of losses or wins? Essentially all of the positive P&L comes from the long run of wins at the end, so is there some repeatable element in that set of trades? The answers to all of these questions may well be no. It is possible that these runs are just the normal expectation for the system, and that any tweaking would be counterproductive, but these are the kinds of questions to ask.

Though the trader's performance is not promising, all is not lost. The breakdown of trades by system (A, B, C, and D) in Table 12.1 shows that the B-class trades might have

a stronger edge. These trades, taken alone, made $5,109 on 17 trades, showing a mean trade of $301 with a standard deviation of $515, achieving statistical significance as well. This trader should possibly consider these trades further, and ask if some of the other trades could be eliminated from the playbook.

This is also a vivid illustration of the fact that neither reward/risk ratio nor winning percentage is particularly important when taken alone. Consider the C trades, which have a very impressive win ratio of 81.3 percent, but an average loss more than three times the size of the average win essentially wipes out any edge. Note also that it is not necessary to calculate expected value in this analysis, as the trade means are equivalent to the expected values for each class of trade.

Standardizing for Risk

One complicating factor is that this trader was not consistent in his risk. Though not visible in Table 12.1, risk per trade on the $100,000 account ranged from a low of $199 to a high of $1,790. The nominal trade risk will vary under a fixed percentage plan (see Chapter 9), but this trader was not executing a fixed percentage plan as trade risks ranged from 0.2 percent to 1.8 percent of equity at the time of entry. Standardizing each trade's P&L as a percentage of the amount risked can provide deeper insight into the trader's performance:

$$\%R = \text{Actual dollar P\&L} \div \text{Initial risk}$$

Table 12.2 shows the result of this analysis.

This time, the results paint a very different picture. On average, the trader made 0.3 times the initial risk on every trade, with a standard deviation of 0.8; this is a much better ratio of return to standard deviation than in the first table. The *t*-test shows that this P&L is statistically significant at the 0.001 level, as there would be less than a 0.1 percent chance of seeing a result this extreme due to random chance. There is a very important point here: The trader's net P&L was a combination of both his position sizing choices and whatever actual edge he may have had. Standardizing for risk removes the position sizing effect and reveals that this trader actually *was* trading with a clear statistical edge, even though it was completely obscured by his position sizing decisions. This second analysis also suggests that there probably *is* an edge to the C-class trades, which

TABLE 12.2 Summary Statistics for %R P&L

System	*N*=	Sum	Mean	StDev	AvgWin	AvgLoss	Win%	*p*=
A	29	2.9X	0.1X	0.8X	0.8X	−0.5X	44.8%	0.249
B	17	13.5X	0.8X	0.8X	1.2X	−0.3X	70.6%	0.001
C	16	6.7X	0.4X	0.7X	0.7X	−0.8X	81.3%	0.013
D	6	0.1X	0.0X	0.6X	0.6X	−0.3X	33.3%	0.463
All	68	23.2X	0.3X	0.8X	0.9X	−0.4X	58.8%	0.000

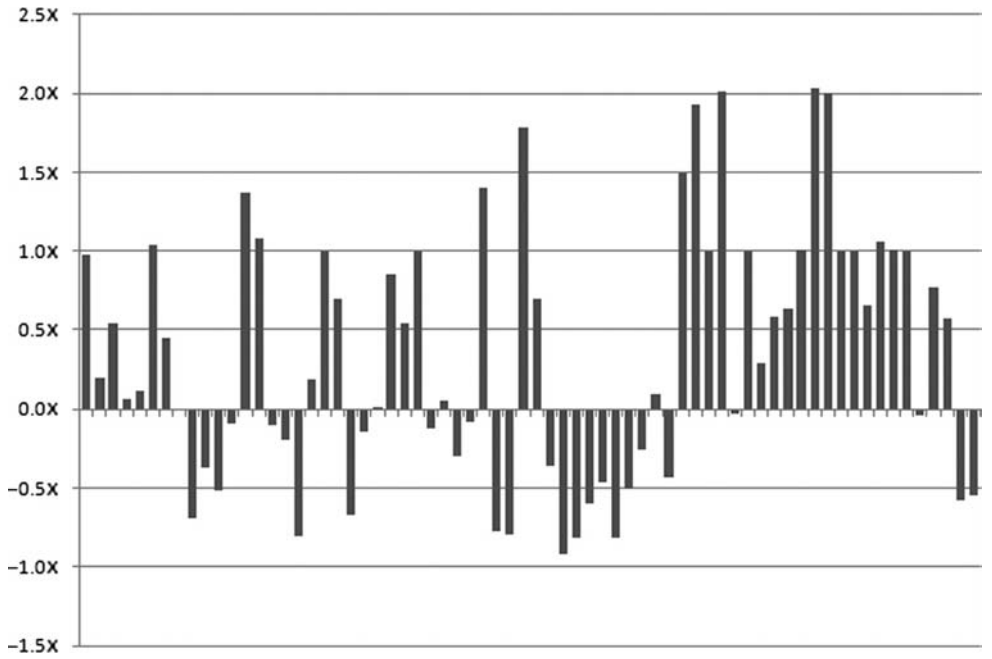

FIGURE 12.3 Trade-by-Trade %R P&L

returned an average of 0.4 times initial risk against a standard deviation of 0.7. Based on this analysis, A-class and D-class trades should be carefully examined, as it appears there may not be an edge.

Compare Figure 12.3, which shows the individual trade results as %R, to Figure 12.2. At first glance they appear to be similar, showing the same runs of wins and losses, but look again. Notice especially that the two large losses that stood out in the middle of Figure 12.2 are normal, insignificant losses when expressed as %R. Is it possible that the trader increased risk on those two trades, taking outsized risks and losses? If so, could those losses have unbalanced him psychologically and led to the series of losing trades? This cannot be answered from a numerical analysis, but these are the kinds of questions that must be asked. Note that, in general, Figure 12.3 is much more consistent—though the trader may have made poor decisions about his position sizing on each trade, it appears that he was working within and respecting his risk limits very well.

This analysis can and should be extended to more detailed investigations of the different classes of trades and different asset classes, if applicable. For now, let's leave this example with one more step. It is possible to generate a pro forma P&L, which assumes that the trader had risked a consistent percentage of account equity on each trade. Figure 12.4 shows this pro forma equity curve, assuming 1 percent fixed fractional risk, against the actual P&L. Note that you cannot simply add the %R results in Table 12.2 to get ending P&L under a fixed fractional scheme because the amount risked is always changing with each trade. In this case, the trader would have ended with $22,935 had he

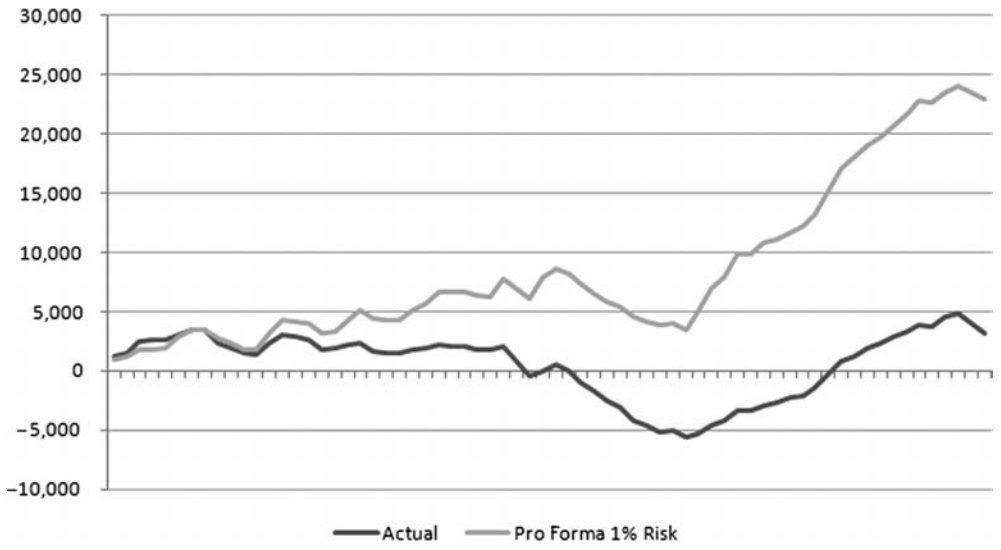

FIGURE 12.4 Cumulative Actual and Pro Forma P&L

simply used 1 percent fixed fractional risk, a stunning gain over his actual $3,158. The path of this pro forma P&L provides another, visual confirmation of what we have seen clearly in the %R analysis—the trader actually did have an edge, but did not apply it very well. The lessons, in this case, are probably simple: perhaps eliminate one or two of the trading systems and consider applying the rest with a fixed fractional risk approach in the next period.

Control Charts

Control charts are tools used in manufacturing to monitor physical processes for early signs of degradation or failure, and are an important quality control tool in diverse fields such as metalworking, electronics, and pharmaceutical manufacturing, among others. They are especially powerful in physical processes that can be tightly controlled, especially with large sample sizes. For instance, an automobile part produced on an assembly line might have tolerances measured in a fraction of a human hair; some variation, within tightly defined limits, is acceptable, but the production line needs a warning when products stray outside these bounds. The same concepts can be applied to trading, with the caveat that trading results are far more variable than manufacturing applications.

There are several different kinds of control charts used in manufacturing applications, but I am going to propose a simplified version that traders may find more immediately applicable. Many of the assumptions of distributions made in classic control charts do not hold in trading situations. For instance, one commonly used chart in manufacturing plots variations on both sides of an average, expecting that the variation should wander more or less freely above and below that average. The quality control manager

takes note when four or five data points appear on one side of the average, and considers this to be an early warning that the process is potentially moving out of the control zone. This degree of control is very unlikely in trading, so the purpose of these charts is slightly different: to present trading results in a graphical format that displays the individual data points, a moving average, and some measure of variation. The first of these charts is the *Standard Deviation Control Chart*, and is produced by the following procedure:

- First, decide what the evaluation period is going to be. Many traders will find weekly analysis to be most applicable, but day traders should do this work on each trading day. (It is also possible to do a trade-by-trade analysis rather than a chronological grouping.)
- Collect a time series of individual (not cumulative) P&Ls for each period.
- Choose a look-back period. The example that follows is done on daily data produced by a day trader, so a look-back period of 20 trading days, which is approximately equal to one trailing calendar month, was chosen. In practice, look-back periods of 10 to 50 are probably most useful.
- For each data point, calculate a moving average of the look-back period as well as the standard deviation. Offset these by one day, so that today you are using yesterday's average and standard deviation. This is important because a large value today will influence the average and standard deviation and, particularly for short look-back periods, will give an understated view of the variation.
- For each data point, plot the raw data point as a bar, the moving average, and bands $+/-N$ times the standard deviation. Remember, it is very unlikely that this data will be normally distributed, so the usual rules of thumb for standard deviation do not apply. Rather, the standard deviation bands are used as a visual representation of the degree of variation in the process. The example that follows uses 2.5 for N, but values of 1.5 to 3.0 may be useful.

Figure 12.5 shows an example of this Standard Deviation Control Chart, based on 190 trading days of a model developing day trader's P&L. Though there is much to be gleaned from this analysis, here are some important points to consider:

A: The trader's variation has been steadily increasing, evidenced by the expanding upper and lower bands. Note that the average (the dark black line in the middle)

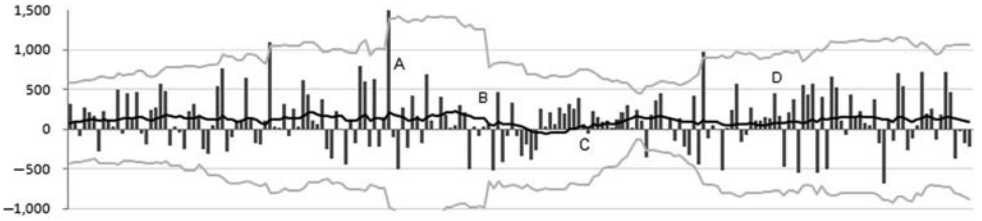

FIGURE 12.5 A Standard Deviation Control Chart

has been consistently above zero, suggesting that the trader is consistently making money. In addition, winning days are much larger than losing days. This is all good.

B: Something happens here. The average takes a turn down, and even crosses to the other side of zero. The trader would have been losing money during this period, but the variation also decreases. It is important to ask the right questions to understand why this happened. What caused the change in the average? Are there behavioral factors at work? Did something shift in the market? Was a new trading system implemented? Perhaps the trader was aware that he was not doing well and reduced trading size. The answers are not on this chart, but the chart can provoke the correct questions.

C: The average recovers after a series of winning days, but there may still be problems. The average is not as consistent, as it turns back down shortly after C, and there are not as many winning days above $500 as at the beginning of the chart.

D: Finally, the trader's P&L seems to be back under control, and the last part of the chart looks very much like the beginning, though the degree of variation (spread of the bands) is higher.

This chart will serve the needs of most traders, but another possibility to consider is the Range Control Chart. This chart, shown in Figure 12.6, can be a useful complement to other P&L analyses. The procedure for creating this chart is simple:

- For each period being evaluated (day, week, or trade), plot the raw data as a bar.
- Choose a look-back period (20 days was used in Figure 12.6), and also plot the high and low for the look-back period.

Consider the lessons of this chart in conjunction with the previous example:

A: Variation is increasing, but it is increasing because of more large winning days. This is good because it suggests that even though the trader's returns may be becoming more volatile, the variation is positive rather than being equally spread around zero.

B: The high/low bands contract, highlighting the reduction of variation already seen on the previous chart.

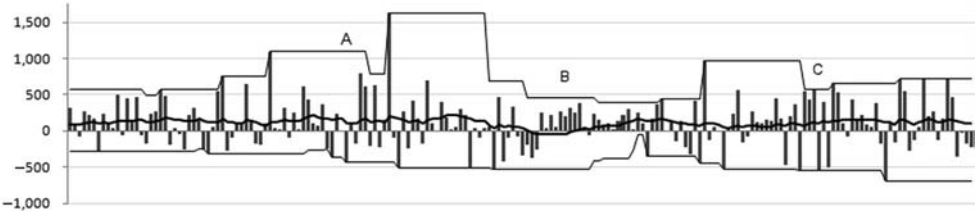

FIGURE 12.6 Range Control Chart

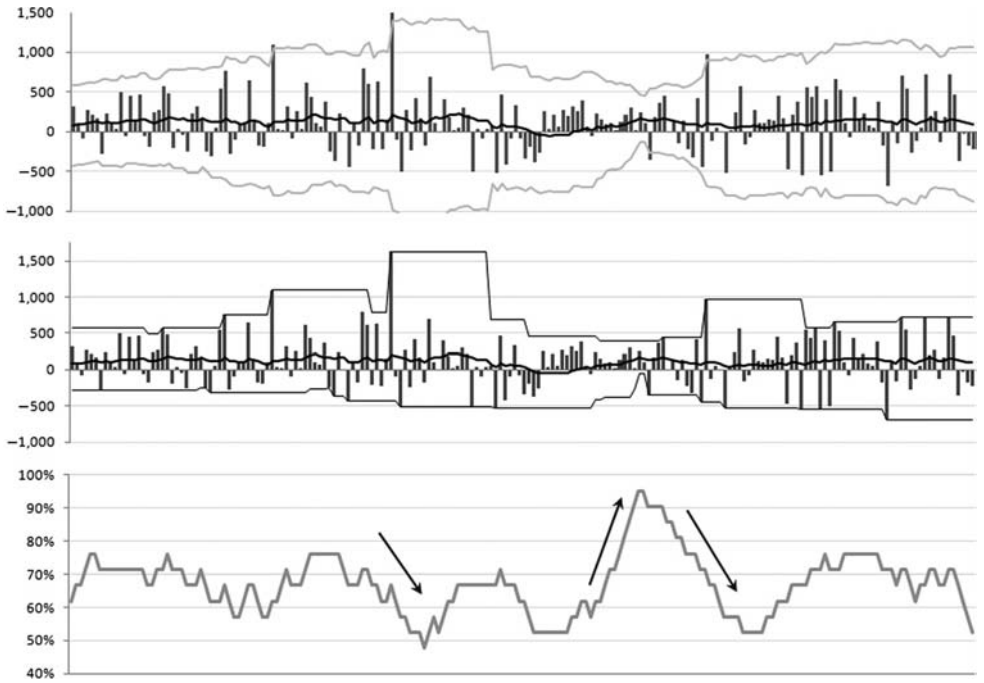

FIGURE 12.7 Three Control Charts

C: Variation is once again increasing, but now more symmetrically. This may temper our optimism at the end of the previous analysis a bit. The trader is making money, but why is more of the volatility to the downside in this period compared to the first part of the chart? Where have the large winning days gone?

One more chart that some traders find useful is a Win Ratio Control Chart, which is calculated by plotting a moving average of the win ratio over a look-back period. Though too much attention on win ratio is counterproductive, a *shift* in win ratio can be a harbinger of trouble to come, and it may be possible to identify problems before they become expensive. Figure 12.7 reproduces the first two charts, with the addition of a Win Ratio Control Chart also using a 20-day look-back period. Consider what information a careful analysis of these three together can give the trader. Note that there are two large declines in win ratio, all of which may point to some potential issues with trading performance. Again, the answers to most trading problems are not on these charts, but in many cases, the right questions are.

SUMMARY

Trading is difficult. Markets are highly random and nearly unpredictable, diverse and unforeseeable risks crop up at every turn, and traders themselves are vulnerable to many

psychological errors and weaknesses. In addition, the process of developing trading skill is a long one, fraught with hardship and obstacles. Many traders fail because they do not have realistic assumptions about the length and challenges of that process, or because they lack some essential requirements for success. This chapter addresses many of those problems, and gives valuable perspective and a much-needed reality check. A trader armed with these, and with the tools and techniques from elsewhere in this book, can develop a plan of attack and move along the path to trading mastery.